

CLAIMS

1. A method for modeling a behavior of an LPAR (logical partition) in a simulated computer operating in a time slice dispatch mode, comprising:
 - beginning a next modeling interval;
 - calculating a resource percentage representing a percentage of total resources allocated to the LPAR;
 - calculating a time slice percentage for the LPAR based on the resource percentage;
 - determining a CP (central processor) percentage representing a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR;
 - and
 - if the CP percentage is greater than the time slice percentage, causing the simulated computer not to dispatch CPs to the LPAR.
2. The method of claim 1, including the further step of repeating each of the recited steps for a next modeling interval.
3. The method of claim 1, wherein the resource percentage is equal to: $100\% - \text{a percentage of resources allocated to all other LPARs running in the simulated computer.}$

4. The method of claim 3, wherein the percentage of resources allocated to all other LPARs is based on a weighting factor specified for each LPAR, a number of logical CPs allocated to each LPAR, and a MIPs (million instructions per second) value for each LPAR.

5. The method of claim 4, wherein the MIPs value represents a maximum consumption that each LPAR could consume in an unrestrained processor.

6. The method of claim 1, wherein:

time slice percentage = $\frac{(\text{resource percentage}) \times (\# \text{ of physical CPs})}{(\# \text{ of logical CPs})}$.

7. A tool for simulating operation of a computer having a system for modeling a behavior of an LPAR operating in a time slice dispatch mode, the modeling system comprising:

a system for calculating a resource percentage, wherein the resource percentage represents a percentage of total resources allocated to the LPAR;

a system for calculating a time slice percentage for the LPAR based on the resource percentage;

a system for determining a CP (central processor) percentage, wherein the CP percentage represents a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR; and

a system for determining causing the computer simulation not to dispatch CPs to the LPAR for a current modeling interval if the CP percentage is greater than the time slice percentage.

8. The tool of claim 7, wherein the resource percentage is equal to: 100% - a percentage of resources allocated to all other LPARs running in the computer simulation.

9. The tool of claim 8, wherein the percentage of resources allocated to all other LPARs is based on a weighting factor specified for each LPAR, a number of logical CPs allocated to each LPAR, and a MIPs (million instructions per second) value for each LPAR.

10. The tool of claim 9, wherein the MIPs value represents a maximum consumption that each LPAR could consume in an unrestrained processor.

11. The tool of claim 7, wherein:

$$\text{time slice percentage} = \frac{(\text{resource percentage}) \times (\# \text{ of physical CPs})}{(\# \text{ of logical CPs})}$$

12. A program product stored on a recordable medium for modeling a behavior of an LPAR (logical partition) in a simulated computer operating in a time slice dispatch mode, comprising:

means for calculating a resource percentage, wherein the resource percentage represents a percentage of total resources allocated to the LPAR;

means for calculating a time slice percentage for the LPAR based on the resource percentage;

means for determining a CP (central processor) percentage, wherein the CP percentage represents a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR; and

means for determining causing the computer simulation not to dispatch CPs to the LPAR for a current modeling interval if the CP percentage is greater than the time slice percentage.

13. The program product of claim 12, wherein the resource percentage is equal to: 100%

- a percentage of resources allocated to all other LPARs.

14. The program product of claim 13, wherein the percentage of resources allocated to all other LPARs is based on a weighting factor specified for each LPAR, a number of logical CPs allocated to each LPAR, and a MIPs (million instructions per second) value for each LPAR.

15. The program product of claim 12, wherein:

time slice percentage = (resource percentage) x (# of physical CPs)

(# of logical CPs).

16. A method for modeling workload performance of a plurality of LPARs (logical partitions) in a computer simulation, comprising:

providing a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs;

setting an initial defined consumption for each model;

running each model and determining an observed consumption for each model;

comparing the observed consumption with the defined consumption for all of the models; and

for each model that has an observed consumption that does not agree with the defined consumption, feeding the observed consumption back to the other models;

adjusting the defined consumption of each model based on the feedback; and

iteratively repeating the running, comparing, feeding and adjusting steps until the observed consumption agrees with the defined consumption for each model.

17. The method of claim 16, wherein the consumption is a measure of processor resources consumed by each LPAR.

18. A computer simulation tool for modeling workload performance of a plurality of LPARs (logical partitions), comprising:

a system for building a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs;

a system for running each model and determining an observed consumption for each model;

a system for comparing the observed consumption with the defined consumption for all of the models;

a system for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption;

a system for adjusting the defined consumption of each model based on the observed consumption feedback; and

a system for iteratively rerunning each model until the observed consumption agrees with the defined consumption for each model.

19. A program product stored on a recordable medium for modeling workload

performance of a plurality of LPARs (logical partitions), comprising:

means for specifying a model for each of a plurality of LPARs, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs;

means for running each model and determining an observed consumption for each model;

means for comparing the observed consumption with the defined consumption for all of the models;

means for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption;

means for adjusting the defined consumption of each model based on the observed consumption feedback; and

means for iteratively rerunning each model until the observed consumption agrees with the defined consumption for each model.

20. A computer simulation tool for modeling LPAR behavior, comprising:
 - a first algorithm for modeling the behavior of an LPAR (logical partition) operating in a time slice dispatch mode; and
 - a second algorithm for modeling the behavior of a plurality of LPARs.
21. The computer simulation tool of claim 15, wherein the first algorithm includes:
 - means for calculating a resource percentage, wherein the resource percentage represents a percentage of total resources allocated to the LPAR;
 - means for calculating a time slice percentage for the LPAR based on the resource percentage;
 - means for determining a CP (central processor) percentage, wherein the CP percentage represents a percentage of time that all physical CPs in a computer being modeled have been allocated to the LPAR; and
 - means for determining causing the computer simulation not to dispatch CPs to the LPAR for a current modeling interval if the CP percentage is greater than the time slice percentage.

22. The computer simulation tool of claim 20, wherein the second algorithm includes:

means for specifying a model for each of a plurality of LPARs, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs;

means for running each model and determining an observed consumption for each model;

means for comparing the observed consumption with the defined consumption for all of the models;

means for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption;

means for adjusting the defined consumption of each model based on the observed consumption feedback; and

means for iteratively rerunning each model until the observed consumption agrees with the defined consumption for each model.